

# Curation is a False God

Where are the alternative proposals for fighting fake news?

## The Call to Arms

Information is a global resource that shapes reality. It's no less important -- and no less valuable -- than oil or land or water. It should come as no surprise it's being fought over.

The weapon of choice on all sides is fake news, an insidious and highly dangerous weapon because it's an invasive species. It won't *kill* real news, it will just out-compete and out-populate it. The nets we cast for news will continue to bring up news, but over time the real-news items will become few among many, and will be camouflaged in the net. What happens when we can't tell the real news from the fake anymore? Your guess is as good as mine.

But *this*, I'm sure of: fake news is not a fringe phenomenon, or a limited-time offering that comes and goes, like the tide, or the McRib. We're allowing people to print unlimited amounts of counterfeit information, and they're getting better at it every day. Why would they stop?

This information economy is doomed. We're chasing our tails trying to solve the unsolvable *counterfeit* part of the problem. We need to focus on solving the *unlimited-amounts* part of the problem.

Fake news is coming for us. How will we defend? We need to defend.

Official Battle Cry: **Curation is not a defense!**

## Curation Makes it Worse

Nothing is qualified to curate truth because nothing has access to the truth. This should be obvious on its face, but especially obvious to anyone who believes in the power of the right to speak one's mind -- the very foundation of the Enlightenment. Curation is a false god. All it does is splinter us infinitely into worlds of delusion where everyone is precisely the flavor of idiot we are, and we all feel high and mighty about it. Curation is the weapon the enemy is fooling us into using against ourselves.

## Where's your head at?

If you got this survey asking you what you think would eventually happen with fake news, which boxes would you check? (Choose all that apply.)

- ☐ It will never be good enough to fool me on any large scale.
- ☐ It will never be good enough to fool large tech companies and news organizations.
- ☐ It will level off after a while.
- ☐ It will only ever fool a small and stupid part of the population.
- ☐ It will never corrupt sites like Wikipedia.
- ☐ It'll all work itself out somehow.

I really *am* curious. I can't check any of those boxes. I am 100% convinced I can be fooled and manipulated. Magicians fool me, even with simple tricks. *People* -- myself included -- fool me with simple tricks. Furthermore, my grip on reality is based wholly on the numbers: I feel sane and sure of it because so many people -- all around the world -- tell me they believe almost exactly the same things I do. I can't *personally* verify *any* of it.

If you wanted to attack my sanity you would do it with a divide-and-conquer strategy. You would break the world down into smaller and smaller belief groups, so that any group I might join would be shrinking. This is fake news' poison. Don't you feel it?

Fake news is not just a bunch of dumb, made-up stories for suckers. Done well, it's an undetectable reality distortion, subtly crafted. Lies are not required. It can be done with smoke and mirrors. The wording of the sentence, the framing of the scene; the details left out, the false implications built in. Maybe that sounds innocuous, but it's not. It can move you very far from where you are now, inch-by-inch, unnoticeable to you. Where does it take you? Wherever you lead it. With every click you make, you're leading it right to the easiest possible place to fool you, because you're always hunting down what you *want* to be true. You're not looking for knowledge, you're looking for affirmation. You're so ready to believe, fake news doesn't have to try very hard to convince you that what you *want* to be true *is* true.

If *I* can be fooled, and *you* can be fooled, then all our organizations are vulnerable as well, because this is a strategy that kills the host by attacking the cells.

So, how do we bind our shared truths together in such a way that the bonds can't be weakened by a steady diet of tasty fake-news snacks? Given how our bodies look in 2021, I don't think we can just assume our shared truths are safe.

## Abstract

A serious problem with our Internet is that anyone can publish as much fake news as they like, for free, and anonymously. The percentage of news that is fake news will continue to increase (for the usual reasons, of course: you can make money, get laid, and win wars by fooling people) until at some point we're simply overwhelmed by it. We as individuals don't have the cognitive abilities to filter unlimited, unsourced information, even though today we still think we do.

The sheer volume of fake news can overwhelm our ability to filter it, even if it's clear which news is fake and which isn't. But it *won't* be clear because fake news will personalize itself to us. It will home in on the borderline between what we *believe* to be true and what we *want* to be true. In this space, it can and will manipulate us, no matter how smart we think we are.

We'll get fooled again and again, maybe, at first, only on stupid things that don't matter, but it will crush our confidence anyway. We'll stop trying to tell the difference; we'll hedge our bets so that we won't be wrong again. This will lead ultimately to a low-trust world, which is not a world I want to live in, and I bet you don't either.

Unless you have something else to offer up -- and you'd better not say curation -- the only weapons we have available to us, are: (1) the creation of a limited and privileged subset of information on the Internet via a publishing tax, and (2) a requirement that verifiable reputations be attached to each piece of information in that subset.

This subset effectively creates a bar (in the beer-drinking sense) with an entrance fee and id check. There will definitely be some liars inside the bar -- our bouncers can't effectively filter them out as they enter, so we won't let them even try. But no matter, you must understand what an upgrade this is. We've moved from an unlimited number of anonymous suspected liars to a limited number of known suspected liars. This is a far better place for trust to flourish!

## TIPT at a High-level

TIPT (The Internet Public Trust) is a proposal for a new Internet service that allows anyone to register the hash (SHA-256, say) of any digital artifact. Together with the hash, the user's real name and the date of registration are stored as metadata, and a fee is collected. The user's name is taken from the credit card the registrant pays with. No anonymous registrations allowed. (In other words, the registrant must be track-down-able in the real world.) The data is stored in a non-distributed, but fully transparent, blockchain ledger.

(Is non-distributed and fully transparent an oxymoron? Am I the moron? If this is the fundamental insight of cryptocurrencies, I apologize for being ignorant. The point I'm trying to make is that TIPT must be fully trusted, both in a mathematical sense, and in a person-on-the-street sense, or else it's worthless. I'm specifying non-distributed only to make it fast and cheap to run.)

When an Internet user-agent -- like a browser -- comes to TIPT with a hash in hand, TIPT's response will indicate whether or not the hash has been registered -- and if it has -- who registered it, and when. If the user agent comes calling with a person's name (aka, a publisher id), TIPT will send back a list of all the artifacts that person has registered, and their dates of registration.

**Registration doesn't indicate ownership.** Registration typically indicates a vouching-for: the registrant was willing to pay with her wallet -- and her reputation -- to boost the signal of a particular digital artifact. It's like paying to give a personalized 5-star rating. That's all it is.

But registration can also be used for proof of authenticity. The New York Times, for example, would want to register its articles so that any reader could check to see if the particular text she's reading actually did come from the Times, without having to go to the Times, by which I mean nytimes.com. (Trust must be intrinsic to -- and travel with -- the artifact. The site could go down, or go away completely, but the artifact lives on.)

In a similar vein, the author of a particular work may want that work to be branded with an incontrovertible timestamp. Put another way, TIPT would provide the digital equivalent of carbon dating, to the microsecond.

(Nothing yet on the Internet is, or can be, properly fixed in time! Not in any sort of robust way, anyhow -- no, not even with the awesome archive.org. TIPT fixes this. It starts the clock of digital history. If nothing else, think of the digital archeologists 1000 years from now. They would kill to have our artifacts locked into

a timestamped blockchain. Timestamps outside of one will be fungible hearsay to them because digital information doesn't age.)

## Day-to-day TIPT

You're viewing something on the internet -- some text, a picture, a video. Your app/browser is presenting the content to you in a frame whose presence signals to you that the content claims to be registered with TIPT. Let's say the frame is a thin, intricate and tasteful black-and white pattern. You consume the content as you normally would, and while you do, in the background, your app/browser computes the hash of the content and sends it off to TIPT. If TIPT responds that the hash isn't registered, the frame warns you by turning red. If TIPT responds with the metadata (author and date of registration), the frame turns green and the metadata is made (optionally) visible. If TIPT isn't answering, the frame turns yellow.

# The Prediction

So why would anyone use TIPT? Why would they *pay* to put something on the Internet when they can already do it for free? The short answer is that, in time, they won't be able to do it for free *and* have their content be noticed, because we're headed for a pay-to-publish Internet whether we like it or not. TIPT is a better pay-to-publish solution than the one(s) we will get if we let it happen organically. And we may avoid a lot of social chaos if we do it now.

Here's the reasoning I convinced myself with.

When we come across a new piece of information, the source of that information is important to us because, as we place the information into our model of the world, its landing spot, and its tightness of fit, are critically affected by the source. If some new information has come to us as a third-hand rumor, for example, it may be put off to the side as a curiosity. But if it has come directly from a trusted friend, the information not only goes in tight, its ramifications across the model are chased down and updated. It's this work involved in updating the ramifications that causes us to be very sensitive to source.

Because if we later learn that some accepted piece of information from a highly-trusted source was false, we need to chase down all

the ramifications again. And the more time that's gone by, the more widely the ramifications have spread themselves across the model, and the more deeply they've become enmeshed in it.

Fake news causes us to engage in excessive ramification chasing. We start to ding the trust-scores of all our sources. As our trusted sources dwindle, our models of the world atrophy, which means that our trust in one another does as well.

Fake news is a weapon: it distorts the perception of the world in ways that benefit the publisher. Because publishing is free on the Internet, and because it can be done anonymously, fake news will proliferate and drown out real news over time. Yes, we can curate ourselves islands in the flood of fake news, but our islands will continually shrink, as even our carefully curated sources themselves fall victim to fake news. The fake news they fail to detect -- and then pass on to us -- will force us to kick them off our islands.

As our islands begin to sink, we'll be desperate for anything that floats in the ocean of mistrust around us. And what will be floating will be *advertisements*, simply because they were paid for. That tiny morsel of value in the packaging of the advertisement will make it float, will make it worth looking at, however self-serving the content of its message might be!

Because it was paid for, it can't come in infinite volume. Nor can the advertisements that arrive to attack it tomorrow. And that makes all the difference! *There are resource constraints.*

Advertisements will come to represent a battleground where the truth can get a fair fight, because there you have to pay for your lies, and stand behind them.

And that's how we end up with a pay-to-publish Internet. Eventually, there will be so much fake news out there that the only way to limit it — without curating it, because curation doesn't work — will be to charge for publishing it. While we'll always *literally* be able to publish on the Internet for free like we do today, there will come a time when we'll have to additionally pay for some number of ad impressions that bait our blog post if we want it to be noticed at all. That will effectively be pay-to-publish, but at some later point, we'll go the rest of the way.

Sites like Facebook will contain nothing but *paid-for* articles — aka, advertisements — with the Times *paying* Facebook to display its articles, or link to them. Facebook (or Google, or whichever company wins this competition) will act as the source of trust (that TIPT should be!), *guaranteeing* that the content it's

displaying has been paid for, by this or that person or organization. A pretty sweet business model.

I'm not clamoring a pay-to-publish Internet, where money decides which information counts, I'm saying it's inevitable. Only a market can filter this stream.